arXiv:1407.7921v3 [math.OC] 8 Jul 2015

Distributed event-triggered coordination for average consensus on weight-balanced digraphs

Cameron Nowzari [a] Jorge Cortés [b]

[a] *Department of Electrical and Systems Engineering, University of Pennsylvania, Philadelphia, PA, 19104, USA*

[b] *Department of Mechanical and Aerospace Engineering, University of California, San Diego, CA, 92093, USA*

Abstract

This paper proposes a novel distributed event-triggered algorithmic solution to the multi-agent average consensus problem for networks whose communication topology is described by weight-balanced, strongly connected digraphs. The proposed event-triggered communication and control strategy does not rely on individual agents having continuous or periodic access to information about the state of their neighbors. In addition, it does not require the agents to have a priori knowledge of any global parameter to execute the algorithm. We show that, under the proposed law, events cannot be triggered an infinite number of times in any finite period (i.e., no Zeno behavior), and that the resulting network executions provably converge to the average of the initial agents' states exponentially fast. We also provide weaker conditions on connectivity under which convergence is guaranteed when the communication topology is switching. Finally, we also propose and analyze a periodic implementation of our algorithm where the relevant triggering functions do not need to be evaluated continuously. Simulations illustrate our results and provide comparisons with other existing algorithms.

Key words: discrete event systems, event-triggered control, average consensus, multi-agent systems, weight-balanced digraphs

1 Introduction

This paper studies the multi-agent average consensus problem, where a group of agents seek to agree on the average of their initial states. Due to its numerous applications in networked systems, many algorithmic solutions exist to this problem; however, a majority of them rely on agents having continuous or periodic availability of information from other agents. Unfortunately, this assumption leads to inefficient implementations in terms of energy consumption, communication bandwidth, congestion, and processor usage. Motivated by these observations, our main goal here is the design of a provably correct distributed event-triggered strategy that prescribes when communication and control updates should occur so that the resulting asynchronous network executions still achieve average consensus.

Literature review: Triggered control seeks to understand the trade-offs between computation, communication, sensing, and actuator effort in achieving a desired task with a guaranteed level of performance. Early works (Åström and Bernhardsson., 2002) consider tuning controller executions to the state evolution of a given system, but the ideas have since then been extended to consider other tasks, see (Heemels et al., 2012) and references therein for a recent overview. Among the many references in the context of multi-agent systems, (Mazo Jr. and Tabuada, 2011) specifies the responsibility of each agent in updating the control signals, (Wang and Lemmon, 2011) considers network scenarios with disturbances, communication delays, and packet drops, and (Stöker et al., 2013) studies decentralized event-based control that incorporates estimators of the interconnection signals among agents. Several works have explored the application of event-triggered ideas to the acquisition of information by the agents. To this end, Xie et al. (2009); Heemels and Donkers (2013); Meng and Chen (2013) combine event-triggered controller updates with sampled data that allows for the periodic evaluation of the triggers. Zhong and Cassandras (2010) drop the need for periodic access to information by considering event-based broadcasts, where agents decide with local information only when to obtain further information about neighbors. Self-triggered control (Anta and Tabuada, 2010; Wang and Lemmon, 2009) relaxes the need for

⋆ A preliminary version was presented as (Nowzari and Cortés, 2014) at the 2014 American Control Conference.

Email addresses: cnowzari@seas.upenn.edu (Cameron Nowzari), cortes@ucsd.edu (Jorge Cortés).

local information by deciding when a future sample of the state should be taken based on the available information from the last sampled state. Team-triggered coordination (Nowzari and Cortés, 2016) combines the strengths of event- and self-triggered control into a unified approach for networked systems.

The literature on multi-agent average consensus is vast, see e.g., (Olfati-Saber et al., 2007; Ren and Beard, 2008; Mesbahi and Egerstedt, 2010) and references therein. Olfati-Saber and Murray (2004) introduce a continuous-time algorithm that achieves asymptotic convergence to average consensus for both undirected and weight-balanced directed graphs. Dimarogonas et al. (2012) build on this algorithm to propose a Lyapunov-based event-triggered strategy that dictates when agents should update their control signals but its implementation relies on each agent having perfect information about their neighbors at all times. The work (Seybotha et al., 2013) uses event-triggered broadcasting with time-dependent triggering functions to provide an algorithm where each agent only requires exact information about itself, rather than its neighbors. However, its implementation requires knowledge of the algebraic connectivity of the network. In addition, the strictly time-dependent nature of the thresholds makes the network executions decoupled from the actual state of the agents. Closer to our treatment here, Garcia et al. (2013) propose an event-triggered broadcasting law with state-dependent triggering functions where agents do not rely on the availability of continuous information about their neighbors (under the assumption that all agents have initial access to a common parameter). This algorithm works for networks with undirected communication topologies and guarantees that all inter-event times are strictly positive, but does not discard the possibility of an infinite number of events happening in a finite time period. We consider here a more general class of communication topologies described by weight-balanced, directed graphs. The works (Gharesifard and Cortés, 2012; Rikos et al., 2014) present provably correct distributed strategies that, given a directed communication topology, allow a network of agents to find such weight edge assignments.

Statement of contributions: Our main contribution is the design and analysis of novel event-triggered broadcasting and controller update strategies to solve the multi-agent average consensus problem over weight-balanced digraphs. With respect to the conference version of this work (Nowzari and Cortés, 2014), the present manuscript introduces new trigger designs, extends the treatment from undirected graphs to weight-balanced digraphs, and provides a comprehensive technical treatment. Our proposed law does not require individual agents to have continuous access to information about the state of their neighbors and is fully distributed in the sense that it does not require any a priori knowledge by agents of global network parameters to execute the algorithm. Our Lyapunov-based design builds on the evolution of the network disagreement to synthesize triggers that agents can evaluate using locally available information to make decisions about when to broadcast their current state to neighbors. In our design, we carefully take into account the discontinuities in the information available to the agents caused by broadcasts received from neighbors and their effect on the feasibility of the resulting implementation. Our analysis shows that the resulting asynchronous network executions are free from Zeno behavior, i.e., only a finite number of events are triggered in any finite time period, and exponentially converge to agreement on the average of all agents' initial states over weight-balanced, strongly connected digraphs. We also provide a lower bound on the exponential convergence rate and characterize the asymptotic convergence of the network under switching topologies that remain weight-balanced and are jointly strongly connected. Lastly, we propose a periodic implementation of our event-triggered design that has agents check the triggers periodically and characterize the sampling period that guarantees correctness. Various simulations illustrate our results.

2 Preliminaries

This section introduces some notational conventions and notions on graph theory. Let $\mathbb{R}$, $\mathbb{R}_{>0}$, $\mathbb{R}_{\geq 0}$, and $\mathbb{Z}_{>0}$ denote the set of real, positive real, nonnegative real, and positive integer numbers, respectively. We denote by $\mathbf{1}_N$ and $\mathbf{0}_N \in \mathbb{R}^N$ the column vectors with entries all equal to one and zero, respectively. We let $\|\cdot\|$ denote the Euclidean norm on $\mathbb{R}^N$. We let $\operatorname{diag}\left(\mathbb{R}^N\right) = \{x \in \mathbb{R}^N \mid x_1 = \cdots = x_N\}$. For a finite set S, we let $|S|$ denote its cardinality. Given $x, y \in \mathbb{R}$, Young's inequality (Hardy et al., 1952) states that, for any $\varepsilon \in \mathbb{R}_{>0}$,

$$xy \leq \frac{x^2}{2\varepsilon} + \frac{\varepsilon y^2}{2}. \tag{1}$$

A weighted directed graph (or weighted digraph) $\mathcal{G} = (V, E, W)$ is comprised of a set of vertices $V = \{1, \ldots, N\}$, directed edges $E \subset V \times V$ and weighted adjacency matrix $W \in \mathbb{R}_{\geq 0}^{N \times N}$. Given an edge $(i, j) \in E$, we refer to j as an out-neighbor of i and i as an in-neighbor of j. The sets of out- and in-neighbors of a given node i are $\mathcal{N}_i^{\text{out}}$ and $\mathcal{N}_i^{\text{in}}$, respectively. The weighted adjacency matrix $W \in \mathbb{R}^{N \times N}$ satisfies $w_{ij} > 0$ if $(i, j) \in E$ and $w_{ij} = 0$ otherwise. A path from vertex i to j is an ordered sequence of vertices such that each intermediate pair of vertices is an edge. A digraph $\mathcal{G}$ is strongly connected if there exists a path from all $i \in V$ to all $j \in V$. The out- and in-degree matrices D^{out} and D^{in} are diagonal matrices where

$$d_i^{\text{out}} = \sum_{j \in \mathcal{N}_i^{\text{out}}} w_{ij}, \quad d_i^{\text{in}} = \sum_{j \in \mathcal{N}_i^{\text{in}}} w_{ji},$$

respectively. A digraph is weight-balanced if $D^{\text{out}} = D^{\text{in}}$. The (weighted) Laplacian matrix is $L = D^{\text{out}} - W$. Based on the structure of L, at least one of its eigenvalues is zero and the rest of them have nonnegative real parts. If the digraph $\mathcal{G}$ is strongly connected, 0 is simple with associated eigenvector $\mathbf{1}_N$. The digraph $\mathcal{G}$ is weight-balanced if and only if $\mathbf{1}_N^T L = \mathbf{0}_N$ if and only if $L_s = \frac{1}{2}(L + L^T)$ is positive semidefinite. For a strongly connected and weight-balanced digraph, zero is a simple eigenvalue of L_s. In this case, we order its eigenvalues as $\lambda_1 = 0 < \lambda_2 \leq \cdots \leq \lambda_N$, and note the inequality

$$x^T L x \geq \lambda_2(L_s)\|x - \frac{1}{N}(\mathbf{1}_N^T x)\mathbf{1}_N\|^2, \tag{2}$$

for all $x \in \mathbb{R}^N$. The following property will also be of use later,

$$\lambda_2(L_s)x^T L x \leq x^T L_s^2 x \leq \lambda_N(L_s)x^T L x. \tag{3}$$

This can be seen by noting that L_s is diagonalizable and rewriting $L_s = S^{-1}DS$, where D is a diagonal matrix containing the eigenvalues of L_s.

3 Problem statement

We consider the multi-agent average consensus problem for a network of N agents. We let $\mathcal{G}$ denote the weight-balanced, strongly connected digraph describing the communication topology of the network. Without loss of generality, we use the convention that an agent i is able to receive information from neighbors in $\mathcal{N}_i^{\text{out}}$ and send information to neighbors in $\mathcal{N}_i^{\text{in}}$. We denote by $x_i \in \mathbb{R}$ the state of agent $i \in \{1, \ldots, N\}$. We consider single-integrator dynamics

$$\dot{x}_i(t) = u_i(t), \tag{4}$$

for all $i \in \{1, \ldots, N\}$. It is well known (Olfati-Saber and Murray, 2004) that the distributed continuous control law

$$u_i(t) = -\sum_{j \in \mathcal{N}_i^{\text{out}}} w_{ij}(x_i(t) - x_j(t)), \tag{5}$$

drives each agent of the system to asymptotically converge to the average of the agents' initial conditions. In compact form, this can be expressed by

$$\dot{x}(t) = -Lx(t),$$

where $x(t) = (x_1(t), \ldots, x_N(t))$ is the column vector of all agent states and L is the Laplacian of $\mathcal{G}$. However, in order to be implemented, this control law requires each agent to continuously access state information about its neighbors and continuously update its control law. Here, we are interested in controller implementations that relax both of these requirements by having agents decide in an opportunistic fashion when to perform these actions. Under this framework, neighbors of a given agent only receive state information from it when this agent decides to broadcast its state to them. Equipped with this information, the neighbors update their respective control laws. We denote by $\widehat{x}_i(t)$ the last broadcast state of agent $i \in \{1, \ldots, N\}$ at any given time $t \in \mathbb{R}_{\geq 0}$. We assume that each agent has continuous access to its own state. We then utilize an event-triggered implementation of the controller (5) given by

$$u_i(t) = -\sum_{j \in \mathcal{N}_i^{\text{out}}} w_{ij}(\widehat{x}_i(t) - \widehat{x}_j(t)). \tag{6}$$

Letting $u(t) = (u_1(t), \ldots, u_N(t)) \in \mathbb{R}^N$ and $\widehat{x} = (\widehat{x}_1, \ldots, \widehat{x}_N) \in \mathbb{R}^N$, we write (6) as

$$u(t) = -L\widehat{x}.$$

Note that although agent i has access to its own state $x_i(t)$, the controller (6) uses the last broadcast state $\widehat{x}_i(t)$. This is to ensure that the average of the agents' initial states is preserved throughout the evolution of the system. More specifically, using this controller, one has

$$\frac{d}{dt}(\mathbf{1}_N^T x(t)) = \mathbf{1}_N^T \dot{x}(t) = -\mathbf{1}_N^T L\widehat{x}(t) = 0, \tag{7}$$

where we have used the fact that $\mathcal{G}$ is weight-balanced.

Our aim is to identify triggers that prescribe in an opportunistic fashion when agents should broadcast their state to their neighbors so that the network converges to the average of the initial agents' states. Given that the average is conserved by (6), all the triggers should enforce is that the agents' states ultimately agree.

4 Distributed trigger design

In this section we synthesize a distributed triggering strategy that prescribes when agents should broadcast state information and update their control signals. Our design builds on the analysis of the evolution of the network disagreement characterized by the following candidate Lyapunov function,

$$V(x) = \frac{1}{2}(x - \bar{x})^T(x - \bar{x}), \tag{8}$$

where $\bar{x} = \frac{1}{N}(\mathbf{1}_N^T x)\mathbf{1}_N$ corresponds to agreement at the average of the states of all agents. The next result characterizes a local condition for all agents in the network such that this candidate Lyapunov function is monotonically nonincreasing.

Proposition 4.1 (Evolution of network disagreement) *For $i \in \{1, \ldots, N\}$, let $a_i \in \mathbb{R}_{>0}$ and denote by*

$e_i(t) = \widehat{x}_i(t) - x_i(t)$ the error between agent i's last broadcast state and its current state at time $t \in \mathbb{R}_{\geq 0}$. Then,

$$\dot{V}(t) \leq -\frac{1}{2}\sum_{i=1}^{N}\sum_{j\in\mathcal{N}_i^{\text{out}}} w_{ij}\left[(1-a_i)(\widehat{x}_i-\widehat{x}_j)^2 - \frac{e_i^2}{a_i}\right].$$

PROOF. Note that, since the average is preserved, cf. (7), under the control law (6), $\bar{x} = \frac{1}{N}(\mathbf{1}_N^T x(0))\mathbf{1}_N$. The function $t \mapsto V(x(t))$ is continuous and piecewise continuously differentiable, with points of discontinuity of $\dot{V}$ corresponding to instants of time where an agent broadcasts its state. Whenever defined, this derivative takes the form

$$\dot{V} = x^T\dot{x} - \bar{x}^T\dot{x} = -x^T L\widehat{x} - \bar{x}^T L\widehat{x} = -x^T L\widehat{x},$$

where we have used that the graph is weight-balanced in the last equality. Let $e = (e_1, \ldots, e_N) \in \mathbb{R}^N$ be the vector of errors of all agents. We can then rewrite $\dot{V}$ as

$$\dot{V} = -\widehat{x}^T L\widehat{x} + e^T L\widehat{x}.$$

Expanding this out yields

$$\dot{V} = -\sum_{i=1}^{N}\sum_{j\in\mathcal{N}_i^{\text{out}}}\left[\frac{1}{2}w_{ij}(\widehat{x}_i-\widehat{x}_j)^2 - e_i w_{ij}(\widehat{x}_i-\widehat{x}_j)\right].$$

Using Young's inequality (1) for each product $e_i(\widehat{x}_i - \widehat{x}_j)$ with $\varepsilon = a_i > 0$ yields,

$$\begin{aligned}\dot{V} &\leq -\sum_{i=1}^{N}\sum_{j\in\mathcal{N}_i^{\text{out}}} w_{ij}\left[\frac{1}{2}(\widehat{x}_i-\widehat{x}_j)^2 - \frac{e_i^2}{2a_i} - \frac{a_i(\widehat{x}_i-\widehat{x}_j)^2}{2}\right] \\ &= -\frac{1}{2}\sum_{i=1}^{N}\sum_{j\in\mathcal{N}_i^{\text{out}}} w_{ij}\left[(1-a_i)(\widehat{x}_i-\widehat{x}_j)^2 - \frac{e_i^2}{a_i}\right],\end{aligned}$$

which concludes the proof. □

From Proposition 4.1, a sufficient condition to ensure that the proposed candidate Lyapunov function V is monotonically decreasing is to maintain

$$\sum_{j\in\mathcal{N}_i^{\text{out}}} w_{ij}\left[(1-a_i)(\widehat{x}_i-\widehat{x}_j)^2 - \frac{e_i^2}{a_i}\right] \geq 0,$$

for all $i \in \{1, \ldots, N\}$ at all times. This is accomplished by ensuring

$$e_i^2 \leq \frac{a_i(1-a_i)}{d_i^{\text{out}}}\sum_{j\in\mathcal{N}_i^{\text{out}}} w_{ij}(\widehat{x}_i-\widehat{x}_j)^2,$$

for all $i \in \{1, \ldots, N\}$. The maximum of the function $a_i(1-a_i)$ in the domain $(0, \infty)$ is attained at $a_i = \frac{1}{2}$, so we have each agent select this value to optimize the trigger design. As a consequence of the above discussion, we have the following result.

Corollary 4.2 *For each $i \in \{1, \ldots, N\}$, let $\sigma_i \in (0, 1)$ and define*

$$f_i(e_i) = e_i^2 - \sigma_i\frac{1}{4d_i^{\text{out}}}\sum_{j\in\mathcal{N}_i^{\text{out}}} w_{ij}(\widehat{x}_i-\widehat{x}_j)^2. \qquad (9)$$

If each agent i enforces the condition $f_i(e_i(t)) \leq 0$ at all times, then

$$\dot{V}(t) \leq -\sum_{i=1}^{N}\frac{1-\sigma_i}{4}\sum_{j\in\mathcal{N}_i^{\text{out}}} w_{ij}(\widehat{x}_i-\widehat{x}_j)^2.$$

(Note that the latter quantity is strictly negative for all $\widehat{x} \notin \text{diag}\left(\mathbb{R}^N\right)$ because the graph is strongly connected).

For each $i \in \{1, \ldots, N\}$, we refer to the function f_i defined in Corollary 4.2 as the *triggering function* and to the condition $f_i(e_i) = 0$ as the *trigger*. Note that the design parameter σ_i affects how flexible the trigger is: as the value of σ_i is selected closer to 1, the trigger is enabled less frequently at the cost of agent i contributing less to the decrease of the Lyapunov function.

An important observation is that, since the triggering function f_i depends on the last broadcast states $\widehat{x}$, a broadcast from a neighbor of i might cause a discontinuity in the evaluation of $f(e_i)$, where just before the update was received, $f_i(e_i) < 0$, and immediately after, $f_i(e_i) > 0$. Such event would make agent i miss the trigger. Thus, rather than prescribing agent $i \in \{1, \ldots, N\}$ to broadcast its state when $f_i(e_i) = 0$, we instead define an event by either

$$f_i(e_i) > 0, \qquad (10)$$

or

$$f_i(e_i) = 0 \quad \text{and} \quad \phi_i \neq 0, \qquad (11)$$

where for convenience, we use the shorthand notation

$$\phi_i = \sum_{j\in\mathcal{N}_i^{\text{out}}} w_{ij}(\widehat{x}_i-\widehat{x}_j)^2 \in \mathbb{R}_{\geq 0}.$$

We note the useful equality $\sum_{i=1}^{N}\phi_i = 2\,\widehat{x}^T L\widehat{x}$. The reasoning behind these triggers is the following. The inequality (10) makes sure that the discontinuities of ϕ_i do not make the agent miss an event. The trigger (11) makes sure that the agent is not required to continuously

broadcast its state to neighbors when its last broadcast state is in agreement with the states received from them.

The triggers (10) and (11) are a generalization of the ones proposed in (Garcia et al., 2013). However, it is unknown whether they are sufficient to exclude the possibility of Zeno behavior in the resulting executions. To address this issue, we prescribe the following additional trigger. Let t^i_{last} be the last time at which agent i broadcast its information to its neighbors $\mathcal{N}_i^{\text{in}}$. If at some time $t \geq t^i_{\text{last}}$, agent i receives new information from a neighbor $j \in \mathcal{N}_i^{\text{out}}$, then i immediately broadcasts its state if

$$t \in (t^i_{\text{last}}, t^i_{\text{last}} + \varepsilon_i). \tag{12}$$

Here, $\varepsilon_i \in \mathbb{R}_{>0}$ is a design parameter selected so that

$$\varepsilon_i < \sqrt{\frac{\sigma_i}{4d_i^{\text{out}} w_i^{\max} |\mathcal{N}_i^{\text{out}}|}}, \tag{13}$$

where $w_i^{\max} = \max_{j \in \mathcal{N}_i^{\text{out}}} w_{ij}$. Our analysis in Section 5 will expand on the role of this bound and the additional trigger in preventing the occurrence of Zeno behavior.

In conclusion, the triggers (10)-(12) form the basis of the EVENT-TRIGGERED COMMUNICATION AND CONTROL LAW, which is formally presented in Table 1.

At all times t agent $i \in \{1, \ldots, N\}$ performs:
1: **if** $f_i(e_i(t)) > 0$ **or** $(f_i(e_i(t)) = 0$ and $\phi_i(t) \neq 0)$ **then**
2: broadcast state information $x_i(t)$ and update control signal
3: **end if**
4: **if** new information $x_j(t)$ is received from some neighbor(s) $j \in \mathcal{N}_i^{\text{out}}$ **then**
5: **if** agent i has broadcast its state at any time $t' \in (t - \varepsilon_i, t)$ **then**
6: broadcast state information $x_i(t)$
7: **end if**
8: update control signal
9: **end if**

Table 1
EVENT-TRIGGERED COMMUNICATION AND CONTROL LAW.

Each time an event is triggered by an agent, say $i \in \{1, \ldots, N\}$, that agent broadcasts its current state to its out-neighbors and updates its control signal, while its in-neighbors $j \in \mathcal{N}_i^{\text{in}}$ update their control signal. This is in contrast to other event-triggered designs, see e.g., (Zhongxin and Zengqiang, 2010; Dimarogonas et al., 2012), where events only correspond to updates of control signals because exact information is available to the agents at all times.

5 Analysis of the event-triggered communication and control law

Here we analyze the properties of the control law (6) in conjunction with the EVENT-TRIGGERED COMMUNICATION AND CONTROL LAW of Section 4. Our first result shows that the network executions are guaranteed not to exhibit Zeno behavior. Its proof illustrates the role played by the additional trigger (12) in facilitating the analysis to establish this property.

Proposition 5.1 (No Zeno behavior) *Given the system* (4) *with control law* (6) *executing the* EVENT-TRIGGERED COMMUNICATION AND CONTROL LAW *over a weight-balanced, strongly connected digraph, the agents will not be required to communicate an infinite number of times in any finite time period.*

PROOF. We are interested in showing here that no agent will broadcast its state an infinite number of times in any finite time period. Our first step consists of showing that, if an agent does not receive new information from neighbors, its inter-event times are lower bounded by a positive constant. Assume agent $i \in \{1, \ldots, N\}$ has just broadcast its state at time t_0, and thus $e_i(t_0) = 0$. For $t \geq t_0$, while no new information is received, $\widehat{x}_i(t)$ and $\widehat{x}_j(t)$ remain constant. Given that $\dot{e}_i = -\dot{x}_i$, the evolution of the error is simply

$$e_i(t) = -(t - t_0)\widehat{z}_i, \tag{14}$$

where, for convenience, we use the shorthand notation $\widehat{z}_i = \sum_{j \in \mathcal{N}_i^{\text{out}}} w_{ij}(\widehat{x}_j - \widehat{x}_i)$. Since we are considering the case when no neighbors of i broadcast information, the trigger (12) is irrelevant. We are then interested in finding the time t^* when $f_i(e_i) = 0$ occurs, triggering a broadcast of agent i's state. If $\widehat{z}_i = 0$, no broadcasts will ever happen ($t^* = \infty$) because $e_i(t) = 0$ for all $t \geq t_0$. Hence, consider the case when $\widehat{z}_i \neq 0$, which in turn implies $\phi_i \neq 0$. Using (14), the trigger (11) prescribes a broadcast at the time $t^* \geq t_0$ satisfying

$$(t^* - t_0)^2 \widehat{z}_i^2 - \sigma_i \frac{1}{4d_i^{\text{out}}} \sum_{j \in \mathcal{N}_i^{\text{out}}} w_{ij}(\widehat{x}_i - \widehat{x}_j)^2 = 0,$$

or, equivalently,

$$(t^* - t_0)^2 = \frac{\sigma_i \sum_{j \in \mathcal{N}_i^{\text{out}}} w_{ij}(\widehat{x}_i - \widehat{x}_j)^2}{4d_i^{\text{out}} \Big(\sum_{j \in \mathcal{N}_i^{\text{out}}} w_{ij}(\widehat{x}_i - \widehat{x}_j) \Big)^2}.$$

Using the fact that $(\sum_{k=1}^p y_k)^2 \leq p \sum_{k=1}^p y_k^2$ for any $y_1, \ldots, y_p \in \mathbb{R}$ and $p \in \mathbb{Z}_{>0}$ (which readily follows from the Cauchy-Schwarz inequality), we obtain

$$\begin{aligned} \Big(\sum_{j \in \mathcal{N}_i^{\text{out}}} w_{ij}(\widehat{x}_i - \widehat{x}_j) \Big)^2 &\leq |\mathcal{N}_i^{\text{out}}| \sum_{j \in \mathcal{N}_i^{\text{out}}} w_{ij}^2 (\widehat{x}_i - \widehat{x}_j)^2 \\ &\leq |\mathcal{N}_i^{\text{out}}| w_i^{\max} \sum_{j \in \mathcal{N}_i^{\text{out}}} w_{ij}(\widehat{x}_i - \widehat{x}_j)^2. \end{aligned} \tag{15}$$

Therefore, we can lower bound the inter-event time by

$$\tau_i = t^* - t_0 \geq \sqrt{\frac{\sigma_i}{4d_i^{\text{out}} w_i^{\max} |\mathcal{N}_i^{\text{out}}|}} > 0,$$

(incidentally, this explains our choice in (13)). Our second step builds on this fact to show that messages cannot be sent an infinite number of times between agents in a finite time period. Let time t_0 be the time at which agent i has broadcast its information to neighbors and thus $e_i(t_0) = 0$. If no information is received by time $t_0 + \varepsilon_i < t_0 + \tau_i$, there is no problem since $\varepsilon_i > 0$, so we now consider the case that at least one neighbor of i broadcasts its information at some time $t_1 \in (t_0, t_0+\varepsilon_i)$. In this case it means that at least one neighbor $j \in \mathcal{N}_i^{\text{out}}$ has broadcast new information, thus agent i would also rebroadcast its information at time t_1 due to trigger (12). Let I denote the set of all agents who have broadcast information at time t_1 (we refer to these agents as synchronized). This means that, as long as no agent $k \notin I$ sends new information to any agent in I, the agents in I will not broadcast new information for at least $\min_{j\in I} \tau_j$ seconds, which includes the original agent i. As before, if no new information is received by any agent in I by time $t_1 + \min_{j\in I} \varepsilon_j$ there is no problem, so we now consider the case that at least one agent k sends new information to some agent $j \in I$ at time $t_2 \in (t_1, t_1 + \min_{j\in I} \varepsilon_j)$. By trigger (12), this would require all agents in I to also broadcast their state information at time t_2 and agent k will now be added to the set I. Reasoning repeatedly in this way, the only way for infinite communications to occur in a finite time period is for an infinite number of agents to be added to the set I, which is not possible because there are only a finite number of agents N. □

Remark 5.2 (Conditions for Zeno) We note here that the introduction of the trigger (12) is sufficient to ensure Zeno behavior does not occur but it is an open problem to determine whether it is also necessary. The design in (Garcia et al., 2013, Corollary 2) specifies triggers of a nature similar to (10)-(11) for undirected graphs and guarantees that no agent undergoes an infinite number of updates at any given instant of time, but does not discard the possibility of an infinite number of updates in a finite time period, as Proposition 5.1 does.•

Next, we establish global exponential convergence.

Theorem 5.3 (Exponential convergence to average consensus) *Given the system* (4) *with control law* (6) *executing the* EVENT-TRIGGERED COMMUNICATION AND CONTROL LAW *over a weight-balanced strongly connected digraph, all agents exponentially converge to the average of the initial states, i.e.,* $\lim_{t\to\infty} x(t) = \bar{x}$.

PROOF. By design, we know that the event-triggers (10)-(11) ensure that, cf. Corollary 4.2,

$$\dot{V} \leq \sum_{i=1}^{N} \frac{\sigma_i - 1}{4} \phi_i. \tag{16}$$

We show that convergence is exponential by establishing that the evolution of V towards 0 is exponential. Define $\sigma_{\max} = \max_{i\in\{1,\dots,N\}} \sigma_i$ to further bounding (16) by

$$\dot{V} \leq \frac{\sigma_{\max} - 1}{4} \sum_{i=1}^{N} \phi_i = \frac{\sigma_{\max} - 1}{2} \widehat{x}^T L \widehat{x}.$$

Given this inequality, our next step is to relate the value of $V(x)$ with $\widehat{x}^T L \widehat{x}$. Note that

$$\begin{aligned} V(x) &\leq \frac{1}{2\lambda_2(L_s)} x^T L x = \frac{1}{2\lambda_2(L_s)} (\widehat{x} - e)^T L (\widehat{x} - e) \\ &= \frac{1}{2\lambda_2(L_s)} \big(\widehat{x}^T L \widehat{x} - 2\widehat{x}^T L_s e + e^T L e \big), \end{aligned}$$

where we have used (2) in the inequality. Now,

$$e^T L e \leq \lambda_N(L_s) \|e\|^2 \leq \lambda_N(L_s) \frac{\sigma_{\max}}{2d_{\min}^{\text{out}}} \widehat{x}^T L \widehat{x},$$

where $d_{\min}^{\text{out}} = \min_{i\in\{1,\dots,N\}} d_i^{\text{out}}$ and we have used $f_i(e_i) \leq 0$ in the second inequality. On the other hand,

$$\begin{aligned} |\widehat{x}^T L_s e| &\leq \|L_s \widehat{x}\| \, \|e\| \leq \sqrt{\lambda_N(L_s) \widehat{x}^T L \widehat{x}} \sqrt{\frac{\sigma_{\max}}{2d_{\min}^{\text{out}}} \widehat{x}^T L \widehat{x}} \\ &= \sqrt{\lambda_N(L_s) \frac{\sigma_{\max}}{2d_{\min}^{\text{out}}}} \widehat{x}^T L \widehat{x}, \end{aligned}$$

where we have used (3) in the second inequality. Putting these bounds together, we obtain

$$V(x) \leq \mathcal{A} \, \widehat{x}^T L \widehat{x},$$

with $\mathcal{A} = \frac{1}{2\lambda_2(L_s)} \big(1 + \sqrt{\lambda_N(L_s) \frac{\sigma_{\max}}{2d_{\min}^{\text{out}}}}\big)^2$. Using this expression in the bound for the Lie derivative, we get

$$\dot{V} \leq \frac{\sigma_{\max} - 1}{2} \widehat{x}^T L \widehat{x} \leq \frac{\sigma_{\max} - 1}{2\mathcal{A}} V(x(t)).$$

This, together with the fact that $t \mapsto V(x(t))$ is continuous and piecewise differentiable implies, using the Comparison Lemma, cf. (Khalil, 2002), that $V(x(t)) \leq V(x(0)) \exp(\frac{\sigma_{\max}-1}{2\mathcal{A}} t)$ and hence the exponential convergence of the network trajectories to the average state. □

The Lyapunov function used in the proof of Theorem 5.3 does not depend on the specific network topology. Therefore, when the communication digraph is time-varying,

this function can be used as a common Lyapunov function to establish asymptotic convergence to average consensus. This observation is key to establish the next result, whose proof we omit for reasons of space.

Proposition 5.4 (Convergence under switching topologies) *Let Ξ_N be the set of weight-balanced digraphs over N vertices. Denote the communication digraph at time t by $\mathcal{G}(t)$. Consider the system (4) with control law (6) executing the* EVENT-TRIGGERED COMMUNICATION AND CONTROL LAW *over a switching digraph, where $t \mapsto \mathcal{G}(t) \in \Xi_N$ is piecewise constant and such that there exists an infinite sequence of contiguous, nonempty, uniformly bounded time intervals over which the union of communication graphs is strongly connected. Then, assuming all agents are aware of who its neighbors are at each time and agents broadcast their state if their neighbors change, all agents asymptotically converge to the average of the initial states.*

6 Periodically checked event-triggered coordination

Here we propose an alternative strategy, termed PERIODIC EVENT-TRIGGERED COMMUNICATION AND CONTROL LAW, where agents only evaluate triggers (10) and (11) periodically, instead of continuously. Specifically, given a sampling period $h \in \mathbb{R}_{>0}$, we let $\{t_\ell\}_{\ell \in \mathbb{Z}_{\geq 0}}$, where $t_{\ell+1} = t_\ell + h$, denote the sequence of times at which agents evaluate the decision of whether to broadcast their state to their neighbors. This type of design is more in line with the constraints imposed by real-time implementations, where individual components work at some given frequency, rather than continuously. An inherent and convenient feature of this strategy is the lack of Zeno behavior (since inter-event times are naturally lower bounded by h), making the need for the additional trigger (12) superfluous. The strategy is formally presented in Table 2.

At times $t \in \{0, h, 2h, \dots\}$, agent $i \in \{1, \dots, N\}$ performs:
1: **if** $f_i(e_i(t)) > 0$ **or** ($f_i(e_i(t)) = 0$ and $\phi_i(t) \neq 0$) **then**
2: broadcast state information $x_i(t)$ and update control signal
3: **end if**
4: **if** new information $x_j(t)$ is received from some neighbor(s) $j \in \mathcal{N}_i^{\text{out}}$ **then**
5: update control signal
6: **end if**

Table 2
PERIODIC EVENT-TRIGGERED COMMUNICATION AND CONTROL LAW.

Each time an agent $i \in \{1, \dots, N\}$ broadcasts, this resets the error to zero, $e_i = 0$. However, because triggers are not evaluated continuously, we no longer have the guarantee $f_i(e_i(t)) \leq 0$ at all times t but, instead, have

$$f_i(e_i(t_\ell)) \leq 0, \tag{17}$$

for $\ell \in \mathbb{Z}_{\geq 0}$. The next result provides a condition on h that guarantees the correctness of our design.

Theorem 6.1 (Exponential convergence under periodic event-triggered communication and control law) *Let $h \in \mathbb{R}_{>0}$ be such that*

$$\sigma_{\max} + 4hw_{\max}|\mathcal{N}_{\max}^{\text{out}}| < 1, \tag{18}$$

where $w_{\max} = \max_{i \in \{1,\dots,N\}} w_i^{\max}$ and $|\mathcal{N}_{\max}^{\text{out}}| = \max_{i \in \{1,\dots,N\}} |\mathcal{N}_i^{\text{out}}|$. Then, given the system (4) with control law (6) executing the PERIODIC EVENT-TRIGGERED COMMUNICATION AND CONTROL LAW *over a weight-balanced strongly connected digraph, all agents exponentially converge to the average of the initial states.*

PROOF. Since (17) is only guaranteed at the sampling times under the PERIODIC EVENT-TRIGGERED COMMUNICATION AND CONTROL LAW, we analyze what happens to the Lyapunov function V in between them. For $t \in [t_\ell, t_{\ell+1})$, note that

$$e(t) = e(t_\ell) + (t - t_\ell)L\widehat{x}(t_\ell).$$

Substituting this expression into $\dot{V}(t) = -\widehat{x}^T(t)L\widehat{x}(t) + e^T(t)L\widehat{x}(t)$, we obtain

$$\dot{V}(t) = -\widehat{x}^T(t_\ell)L\widehat{x}(t_\ell) + e^T(t_\ell)L\widehat{x}(t_\ell) + (t - t_\ell)\widehat{x}^T(t_\ell)L^TL\widehat{x}(t_\ell),$$

for all $t \in [t_\ell, t_{\ell+1})$. For a simpler exposition, we drop all arguments referring to time t_ℓ in the sequel. Following the same line of reasoning as in Proposition 4.1 yields

$$\dot{V}(t) \leq \sum_{i=1}^N \frac{\sigma_i - 1}{4}\phi_i + (t - t_\ell)\widehat{x}^TL^TL\widehat{x}.$$

Using (15), we bound

$$\begin{aligned}\widehat{x}^TL^TL\widehat{x} &= \sum_{i=1}^N \Big(\sum_{j \in \mathcal{N}_i^{\text{out}}} w_{ij}(\widehat{x}_i - \widehat{x}_j)\Big)^2 \\ &\leq \sum_{i=1}^N |\mathcal{N}_i^{\text{out}}| w_i^{\max} \sum_{j \in \mathcal{N}_i^{\text{out}}} w_{ij}(\widehat{x}_i - \widehat{x}_j)^2 \\ &= |\mathcal{N}_{\max}^{\text{out}}| w_{\max} \sum_{i=1}^N \phi_i. \end{aligned} \tag{19}$$

Hence, for $t \in [t_\ell, t_{\ell+1})$,

$$\dot{V}(t) \leq \sum_{i=1}^{N} \Big(\frac{\sigma_i - 1}{4} + h w_{\max} |\mathcal{N}_{\max}^{\text{out}}| \Big) \phi_i$$
$$\leq \Big(-\frac{1}{2} + \frac{\sigma_{\max}}{2} + 2h w_{\max} |\mathcal{N}_{\max}^{\text{out}}| \Big) \widehat{x}^T L \widehat{x}.$$

Under (18), a reasoning similar to the proof of Theorem 5.3 using (19) leads to finding $\mathcal{B} > 0$ such that

$$\dot{V}(t) \leq \frac{1}{2\mathcal{B}} \Big(\sigma_{\max} + 4h w_{\max} |\mathcal{N}_{\max}^{\text{out}}| - 1 \Big) V(x(t)),$$

which implies the result. □

7 Simulations

This section illustrates the performance of the proposed algorithms in simulation. Figure 1 shows a comparison of the EVENT-TRIGGERED COMMUNICATION AND CONTROL LAW with the algorithm proposed in (Garcia et al., 2013) for undirected graphs over a network of 5 agents. Both algorithms operate under the dynamics (4) with control law (6), and differ in the way events are triggered. The algorithm in (Garcia et al., 2013) requires all network agents to have knowledge of an a priori chosen common parameter $a \in \mathbb{R}_{>0}$, which we set here to $a = 0.2$. Figure 1(a) shows the evolution of the Lyapunov function V and Figure 1(b) shows the number of events triggered over time by each strategy.

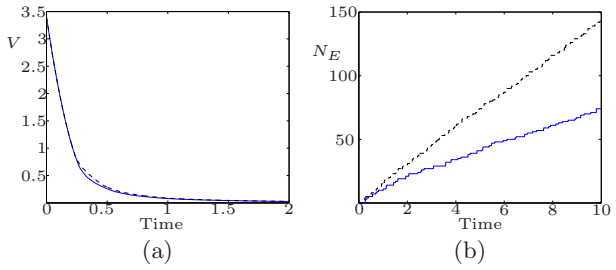

Fig. 1. Plots of (a) the evolution of the Lyapunov function V and (b) the total number N_E of events of the EVENT-TRIGGERED COMMUNICATION AND CONTROL LAW with $\sigma_i = 0.999$ for all i (solid blue) and the algorithm proposed in (Garcia et al., 2013) with $a = 0.2$ (dashed black). The network consists of 5 agents with communication topology described by the undirected graph $(\{1,\dots,5\}, \{(1,2),(1,3),(2,4),(4,5)\})$. The initial condition is $x(0) = [-1,0,2,2,1]^T$.

Figure 2 shows an execution of EVENT-TRIGGERED COMMUNICATION AND CONTROL LAW over a network of 5 agents whose communication topology is described by a weight-balanced digraph. In this case, we do not compare it against the algorithm in Garcia et al. (2013) because the latter is only designed to work for undirected graphs.

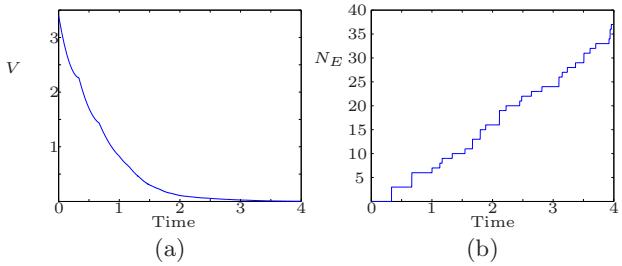

Fig. 2. Plots of (a) the evolution of the Lyapunov function V and (b) the total number N_E of events of the EVENT-TRIGGERED COMMUNICATION AND CONTROL LAW with $\sigma_i = 0.999$ for all i. The network consists of 5 agents with communication topology described by the weight-balanced digraph $(\{1,\dots,5\}, \{(1,2),(2,3),(2,4),(3,4),(4,5),(5,1),(5,2)\})$ with weights $(1,1,0.5,1,1.5,1,0.5)$. The initial condition is $x(0) = [-1,0,2,2,1]^T$.

We have also compared the PERIODIC EVENT-TRIGGERED COMMUNICATION AND CONTROL LAW with a periodic implementation of Laplacian consensus, cf. (Olfati-Saber et al., 2007). For the latter, trajectories are guaranteed to converge if the timestep satisfies $h < 1/d^{\max}$, where $d^{\max}$ is the maximum degree of the graph $\mathcal{G}$. Figure 3 shows this comparison using $h = 0.1$ and also demonstrates the effect of $\{\sigma_i\}_{i=1}^N$ on the executions of the PERIODIC EVENT-TRIGGERED COMMUNICATION AND CONTROL LAW. For simplicity, we have used $\sigma_i = \sigma$ to be the same for all agents in each execution. One can observe the trade-off between communication and convergence rate for varying σ: higher σ results in less communication but slower convergence compared to smaller values of σ.

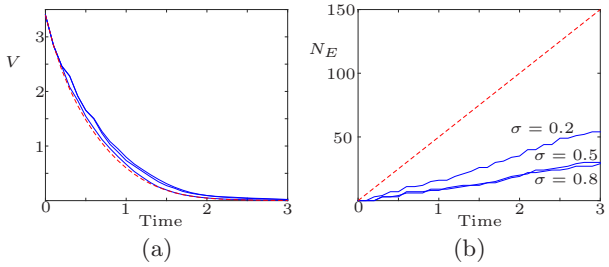

Fig. 3. Plots of (a) the evolution of the Lyapunov function V and (b) the total number N_E of events of the PERIODIC EVENT-TRIGGERED COMMUNICATION AND CONTROL LAW (with varying σ, solid blue) and a standard periodic Laplacian consensus algorithm (dashed red) with timestep $h = 0.1$. Network and initial condition are as in Figure 2.

8 Conclusions

We have proposed novel event-triggered communication and control strategies for the multi-agent average consensus problem. Among the novelties of our first design, we highlight that it works over weight-balanced directed communication topologies, does not require individual agents to continuously access information about

the states of their neighbors, and does not necessitate a priori agent knowledge of global network parameters to execute the algorithm. We have shown that our algorithms exclude the possibility of Zeno behavior and identified conditions such that the network state exponentially converges to agreement on the initial average of the agents' state. We have also provided a lower bound on the convergence rate and characterized the network convergence when the topology is switching under a weaker form of connectivity. Finally, we have developed a periodic implementation of our event-triggered law that relaxes the need for agents to evaluate the relevant triggering functions continuously and provided a sufficient condition on the sampling period that guarantee its the asymptotic correctness. Future work will explore scenarios with more general dynamics and physical sources of error such as communication delays or packet drops, the extension of our design and results to distributed convex optimization and other coordination tasks, and further analysis of trigger designs that rule out the possibility of Zeno behavior.

Acknowledgments

This research was supported in part by NSF award CNS-1329619.

References

A. Anta and P. Tabuada. To sample or not to sample: self-triggered control for nonlinear systems. *IEEE Transactions on Automatic Control*, 55(9):2030–2042, 2010.

K. J. Åström and B. M. Bernhardsson. Comparison of Riemann and Lebesgue sampling for first order stochastic systems. In *IEEE Conf. on Decision and Control*, pages 2011–2016, Las Vegas, NV, December 2002.

D. V. Dimarogonas, E. Frazzoli, and K. H. Johansson. Distributed event-triggered control for multi-agent systems. *IEEE Transactions on Automatic Control*, 57(5):1291–1297, 2012.

E. Garcia, Y. Cao, H. Yu, P. Antsaklis, and D. Casbeer. Decentralised event-triggered cooperative control with limited communication. *International Journal of Control*, 86(9):1479–1488, 2013.

B. Gharesifard and J. Cortés. Distributed strategies for generating weight-balanced and doubly stochastic digraphs. *European Journal of Control*, 18(6):539–557, 2012.

G. H. Hardy, J. E. Littlewood, and G. Polya. *Inequalities*. Cambridge University Press, Cambridge, UK, 1952.

W. P. M. H. Heemels and M. C. F. Donkers. Model-based periodic event-triggered control for linear systems. *Automatica*, 49(3):698–711, 2013.

W. P. M. H. Heemels, K. H. Johansson, and P. Tabuada. An introduction to event-triggered and self-triggered control. In *IEEE Conf. on Decision and Control*, pages 3270–3285, Maui, HI, 2012.

H. K. Khalil. *Nonlinear Systems*. Prentice Hall, 3 edition, 2002. ISBN 0130673897.

M. Mazo Jr. and P. Tabuada. Decentralized event-triggered control over wireless sensor/actuator networks. *IEEE Transactions on Automatic Control*, 56 (10):2456–2461, 2011.

X. Meng and T. Chen. Event based agreement protocols for multi-agent networks. *Automatica*, 49(7): 2125–2132, 2013.

M. Mesbahi and M. Egerstedt. *Graph Theoretic Methods in Multiagent Networks*. Applied Mathematics Series. Princeton University Press, 2010.

C. Nowzari and J. Cortés. Zeno-free, distributed event-triggered communication and control for multi-agent average consensus. In *American Control Conference*, pages 2148–2153, Portland, OR, 2014.

C. Nowzari and J. Cortés. Team-triggered coordination for real-time control of networked cyberphysical systems. *IEEE Transactions on Automatic Control*, 61 (1), 2016. To appear.

R. Olfati-Saber and R. M. Murray. Consensus problems in networks of agents with switching topology and time-delays. *IEEE Transactions on Automatic Control*, 49(9):1520–1533, 2004.

R. Olfati-Saber, J. A. Fax, and R. M. Murray. Consensus and cooperation in networked multi-agent systems. *Proceedings of the IEEE*, 95(1):215–233, 2007.

W. Ren and R. W. Beard. *Distributed Consensus in Multi-Vehicle Cooperative Control*. Communications and Control Engineering. Springer, 2008. ISBN 978-1-84800-014-8.

A. Rikos, T. Charalambous, and C. N. Hadjicostis. Distributed weight balancing over digraphs. *IEEE Transactions on Control of Network Systems*, 2014. To appear.

G. S. Seybotha, D. V. Dimarogonas, and K. H. Johansson. Event-based broadcasting for multi-agent average consensus. *Automatica*, 49(1):245–252, 2013.

C. Stöker, D. Vey, and J. Lunze. Decentralized event-based control: Stability analysis and experimental evaluation. *Nonlinear Analysis: Hybrid Systems*, 10: 141–155, 2013.

X. Wang and M. D. Lemmon. Self-triggered feedback control systems with finite-gain L_2 stability. *IEEE Transactions on Automatic Control*, 54(3):452–467, 2009.

X. Wang and M. D. Lemmon. Event-triggering in distributed networked control systems. *IEEE Transactions on Automatic Control*, 56(3):586–601, 2011.

G. Xie, H. Liu, L. Wang, and Y. Jia. Consensus in networked multi-agent systems via sampled control: fixed topology case. In *American Control Conference*, pages 3902–3907, St. Louis, MO, 2009.

M. Zhong and C. G. Cassandras. Asynchronous distributed optimization with event-driven communication. *IEEE Transactions on Automatic Control*, 55 (12):2735–2750, 2010.

L. Zhongxin and C. Zengqiang. Event-triggered average-consensus for multi-agent systems. In *Chinese Control Conference*, pages 4506–4511, July 2010.